MATTHEW D. POWERS (Bar No. 104795)
matthew.powers@weil.com
EDWARD R. REINES (Bar No. 135960)
edward.reines@weil.com
JEFFREY G. HOMRIG (Bar No. 215890)
jeffrey.homrig@weil.com
JILL J. HO (Bar No. 236349)
jill.ho@weil.com
WEIL, GOTSHAL & MANGES LLP
Silicon Valley Office
201 Redwood Shores Parkway
Redwood Shores, CA 94065
Telephone: (650) 802-3000
Facsimile: (650) 802-3100

Attorneys for Plaintiff-Counterclaim Defendant
NETAPP, INC.,

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

SAN FRANCISCO DIVISION

NETWORK APPLIANCE, INC. Plaintiff-Counterclaim Defendant, v. SUN MICROSYSTEMS, INC. Defendant-Counterclaim Plaintiff.	Case No. C-07-06053 EDL **PLAINTIFF NETAPP, INC.'S OPENING CLAIM CONSTRUCTION BRIEF** Hon. Elizabeth D. Laporte Complaint Filed: September 5, 2007 Hearing Date: August 27, 2008

TABLE OF CONTENTS

Page

INTRODUCTION 1

LEGAL FRAMEWORK 1

ARGUMENT 2

I. U.S. PATENT NO. 5,925,106 2

A. Background 2

B. "Domain Name" (Claims 1, 9, 14, 22, 27, and 35) 5

1. The Specification Supports NetApp's Construction 6

2. The Prosecution History Supports NetApp's Construction 7

3. Extrinsic Evidence Supports NetApp's Construction 7

4. Sun's Proposed Construction is Wrong 9

C. "Server Identification Data" (Claims 1, 4, 9-10, 14, 17-18, 23-24, 27, 30, 35-36) 10

1. The Specification and File History Support NetApp's Construction 10

2. The File History Supports NetApp's Construction 11

3. Sun's Proposed Construction Is Wrong 11

II. U.S. PATENT NO. 5,459,857 13

A. Background 13

1. The Plain Language Of The Claims Demonstrates That The Required Transmission Of The "Written Record" Occurs Only After The Record Is Written To The First Redundancy Group, and That the Required "Selecting" Happens Only After the Data Records are Received 16

2. All 54 Claims Of The '857 Patent Should Not Be Expected To Cover The Most Specific Disclosure Of The Disclosed Embodiment 17

III. U.S. PATENT NO. 5,749,095 18

A. Background 18

B. "Completing [a] Write Operation Within [a] Local Processing Node" / "Completing [a] Write Operation With Respect to [a] Processor" (Claims 1, 11, and 17) 22

IV. U.S. PATENT NO. 6,873,630 26

A. Background 26

B. "Portion [of a] Communication" (Claims 3, 5, 12-15, 21-22, 45, 48, 50, 52-53, 73, 76) 29

1. A "portion of a communication" is a "fraction or portion of a frame" 29

C. "Element [of a] Communication" / "Element [of a] Portion" / "Elements" / " Element of a Portion" (Claims 3, 5, 8, 13-15, 21-23, 45-48, 50, 52, 73, 76, 79, 89, 113, 114, 116, 117) 33

TABLE OF CONTENTS
(continued)

Page

1. An "Element" is a Part of a "Mini-frame" 33
2. An "Element" is Transmitted Across a Single Channel..... 34
3. An "Element" is Individually Encoded For Transmission..... 35

TABLE OF AUTHORITIES

Page(s)

CASES

Adang v. Fischhoff,
286 F.3d 1346 (Fed. Cir. 2002) 16

Apple Computer, Inc., v. Articulate Systems, Inc.,
234 F.3d 14 (Fed. Cir. 2000) 12, 24

AstraZeneca AB, Aktiebolaget Hassle, KBI-E, Inc. v. Mutual Pharmaceutical Co. Inc., 384 F.3d 1333 (Fed. Cir. 2004) 31

BellAtlantic Network Services v. Covad Communications. Group, Inc.,
262 F.3d 1258 (Fed. Cir. 2001) 30

Decisioning.Com, Inc. v. Federated Department Stores, Inc.,
527 F.3d 1300 (Fed. Cir. 2008) 2

Gillespie v. Dywidag Systems International,, USA,
501 F.3d 1285 (Fed. Cir. 2007) 7, 11

Helmsderfer v. Bobrick Washroom Equipment, Inc.,
527 F.3d 1379 (Fed. Cir. 2008) 15, 17

In re Hyatt,
708 F.2d 712 (Fed. Cir. 1983) 16

Markman v. Westview Instruments, Inc.,
52 F.3d 967 (Fed. Cir. 1995), *aff'd* 517 U.S. 370 (1996) 2

Microsoft Corp. v. Multi-Technology Systems, Inc.,
357 F.3d 1340 (Fed. Cir. 2004) 5

On Demand Machine Corp. v. Ingram Industries, Inc.,
442 F.3d 1331 (Fed. Cir. 2006) 31

In re Oppedahl & Larson LLP,
373 F.3d 1171 (Fed. Cir. 2004) 8

Pause Technology LLC v. TiVo Inc.,
419 F.3d 1326 (Fed. Cir. 2005) 16

Phillips v. AWH Corp., 415 F.3d 1303 (Fed. Cir. 2005) 1, 2

Renishaw PLC v. Marposs Societa' per Azioni,
158 F.3d 1243 (Fed. Cir. 1998) 1

Resonate Inc. v. Alteon Websystems, Inc.,
338 F.3d 1360 (Fed. Cir. 2003) 9

TABLE OF AUTHORITIES
(continued)

Page(s)

Sinorgchem Co., Shandong v. International Trade Commission,
511 F.3d 1132 (Fed. Cir. 2007) 17

Teva Pharmaceuticals USA, Inc. v. Pfizer Inc.,
405 F.3d 990 (Fed. Cir. 2005) 8

Verizon Services Corp. v. Vonage Holdings Corp.,
503 F.3d 1295 (Fed. Cir. 2007) 12

STATUTES

15 U.S.C. § 1127 9

MISCELLANEOUS

D. Crouch, *The Rising Size and Complexity of the Legal Document*,
Legal Studies Research Paper Series, Research Paper No. 2008-04,
Univ. of Mo. School of Law (2008) 15

David Dale, *The Road to 10- Gigabit Ethernet*, § 6 (2006) 27

Microsoft Press Computer Dictionary, (3d ed. 1997) 7

Webster's New World Dictionary of Computer Terms (8th ed. 2000) 8

INTRODUCTION

Plaintiff and Counterclaim Defendant NetApp, Inc. submits this Opening Claim Construction Brief pursuant to the Court's February 12, 2008 Case Management Scheduling Order and the Court's June 17, 2008 Order Regarding Claim Construction. There are fourteen terms currently scheduled to be construed by the Court during this claim construction process, each party having proposed seven of these terms. The parties have agreed that each party will submit its Opening Brief on the seven terms that it proposed for construction during this process. NetApp therefore submits this brief in support of its proposed constructions of terms in the following four Sun patents-in-suit: United States Patent Nos. 5,925,106 ("the '106 patent"); 5,459,857 ("the '857 patent"); 5,749,095 ("the '095 patent"); and 6,873,630 ("the '630 patent").

LEGAL FRAMEWORK

The fundamental principles of claim construction are well known to the Court. Accordingly, the following is simply a concise synthesis of the law applied by NetApp in this brief.

"It is a bedrock principle of patent law that the claims of a patent define the invention to which the patentee is entitled the right to exclude." *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312 (Fed. Cir. 2005). However, these "claims must be read in view of the specification, of which they are a part." *Id.* at 1315. "Ultimately, the interpretation to be given a term can only be determined and confirmed with a full understanding of what the inventors actually invented and intended to envelop with the claim. The construction that [1] stays true to the claim language and [2] most naturally aligns with the patent's description of the invention will be, in the end, the correct construction." *Id.* at 1316 (quoting *Renishaw PLC v. Marposs Societa' per Azioni*, 158 F.3d 1243, 1250 (Fed. Cir. 1998)). "[T]here is sometimes a fine line between reading a claim in light of the specification, and reading a limitation into the claim from the specification. [The Federal Circuit] has recognized that attempting to resolve that problem in the context of the particular patent is likely to capture the scope of the actual invention more accurately than either strictly limiting the scope of the claims to the embodiments disclosed in the specification or

divorcing the claim language from the specification, and, thus, that there can be no magic formula or catechism for conducting claim construction." *See Decisioning.Com, Inc. v. Federated Department Stores, Inc.*, 527 F.3d 1300, 1307-1308 (Fed. Cir. 2008) (internal quotations and citations omitted).

In addition, the prosecution history "can often inform the meaning of the claim language by demonstrating how the inventor understood the invention and whether the inventor limited the invention in the course of prosecution, making the claim scope narrower than it otherwise should be." *Phillips*, 415 F.3d at 1317.

Courts may also utilize extrinsic evidence—which "consists of all evidence external to the patent and prosecution history, including expert and inventor testimony, dictionaries, and learned treatises"—to help determine the meaning of disputed claim terms. *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 980 (Fed. Cir. 1995) (*en banc*), *aff'd* 517 U.S. 370 (1996). Extrinsic evidence, however, is generally considered less helpful in determining the legal meaning of claim terms than intrinsic evidence. *Phillips*, 415 F.3d at 1317. Extrinsic evidence can be helpful in providing context for the disputed term and the technology at issue. *Id.* at 1584; *see also Markman*, 52 F.3d at 986.

ARGUMENT

I. U.S. PATENT NO. 5,925,106

A. Background

Sun's '106 Patent is one of those patents that give early Internet patents a bad name. The entirety of the specification is dedicated to describing a facile mechanism for helping a novice user of the Internet better understand the source of a webpage by adding a descriptive reference to something other than the web address itself. The patent proposes to accomplish this by putting the name of the web page owner inside parentheses near the web address so that the novice user will be less likely to become disoriented while "surfing" online. So, for example, the Court's web address (URL), http://cand.uscourts.gov/, reflects that the web pages on the Court website are from the domain "cand.uscourts.gov." To the uninitiated, it might not be clear from the domain name "cand.uscourts.gov" that this website is sponsored by the Northern District of

California. Thus, according to the patent, adding the parenthetical phrase "(Northern District of California)" after the web address on the user's computer screen will help the novice user of the Internet identify the owner of the website. Referred to as "server identification data," this "human-friendly" parenthetical is the essence of the purported invention.

The '106 specification reveals one embodiment alone: a computer system intended to make it easier for inexperienced users of the World Wide Web to understand the origin of the Internet web pages they view with their web browsers. In 1996 the WWW was growing in popularity, but there were still users completely inexperienced with any aspect of computer networks. Use of the web relies heavily upon Internet domain names and Uniform Resource Locators (URLs). Internet domain names are registered in the Internet's registry of domain names and linked to server IP addresses through the Internet's Domain Name System (DNS). URLs are the addresses of a particular format found in the location bar of a web browser, which the web browser uses to access a web page or other resource on the Internet. According to the patent, Internet domain names and URLs were often ambiguous, non-descriptive, or complex as an indicator of the owner of a website, and therefore were intimidating and made it difficult for inexperienced users to understand who owned the website they were viewing. *See* '106 at 1:54-2:3. According to the '106 patent, "WWW users often become disoriented while navigating through the WWW…. [despite that] each WWW page (webpage) is stored on a specific server…. [i]n the prior art, this connection between server and information is not clearly indicated." *See id.* at 1:45-53. The patent further asserts that inexperienced users were therefore unable "to accurately identify the organization that provides the server," and therefore lacked a basic understanding of "the information structure of the WWW." *See id.* at 2:4-7; s*ee also, generally,* Almeroth Decl.

The alleged invention of the '106 patent attempts to overcome this supposed problem by displaying information about the server responsible for a web page, and in particular its owner:

> WWW browser users need to have access to information about a server to better understand the information structure of the WWW. That is, the connection between server and information residing on the server must be clear. Further, the maintainer of a server would like to be able to display data that describes the server to the user

> of a WWW browser. This can be accomplished by displaying, for the example above:
>
> http://www.sun.com/index.html ***(Sun Microsystems, Inc. WWW server)***
>
> With ***the server's identification data appended to the URL***, the connection between the server and the information it provides is clearly indicated. Thus, a WWW user knows more about the informational structure of the WWW.

'106 at 2:7-20.[1]

Consistent with the object of the purported invention, the claimed apparatus and method exploit the already-existing features of HTML, HTTP, existing web browsers, and the Domain Name System ("DNS"), as well as utilities used to access information available through DNS and the Internet's domain name registry, in order to inform inexperienced users about the origin of the information they are viewing.

Figures 3B and 3C illustrate the alleged distinction between the purported invention and the prior art:

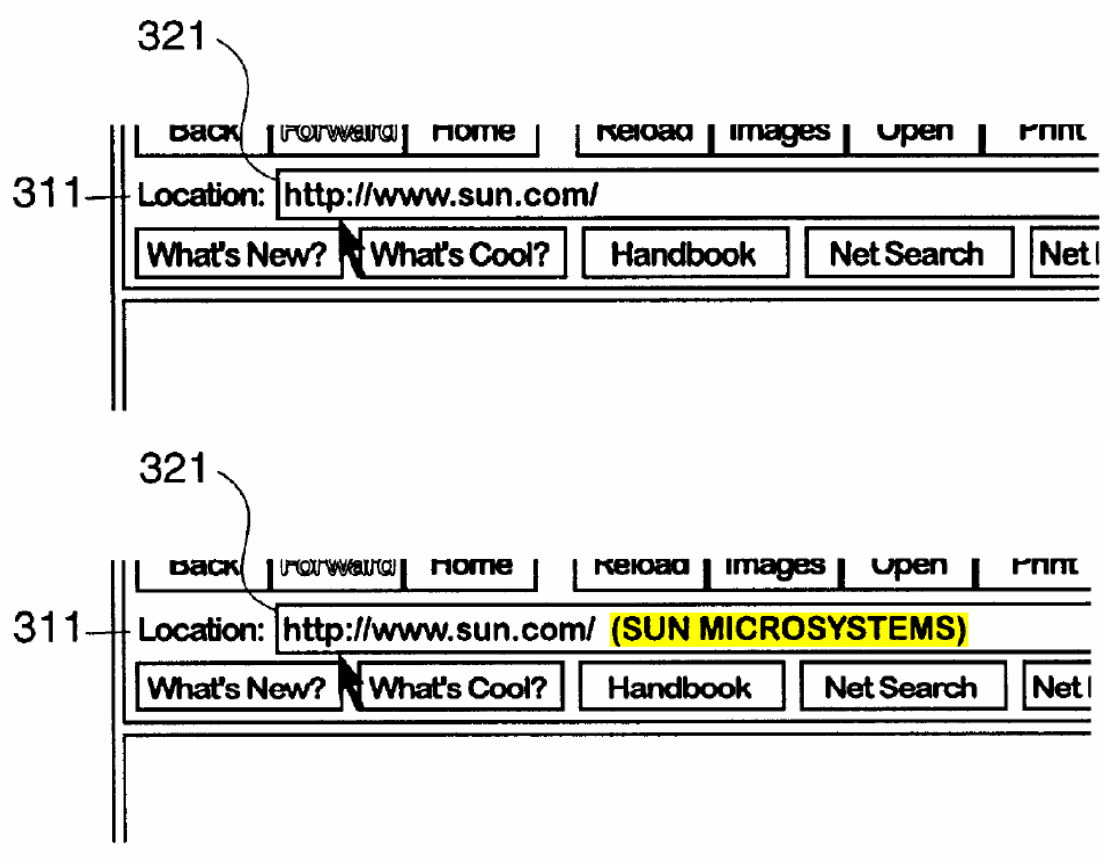

[1] Emphasis added throughout unless otherwise noted.

Like many touted innovations from the early days of the World Wide Web (WWW), the '106 patent dresses up a very simple thought in the trappings of high technology engineering. The patent is already problematic in this respect, enabling only a minor addition to the emerging Internet—if it enables anything at all. Indeed, the claims are broad beyond any rational comparison to the "invention" disclosed in the specification. But now the constructions proposed by Sun would exacerbate the problem by broadening the already broad claims still further. Under Sun's proposed construction, the claims would read on virtually every connection between computers resulting in any output visible to a user. This is therefore a textbook instance of the kind of patent that should be construed to ensure the claims restrict the public use of only the specific invention actually described by the patentee. As the Federal Circuit explained, where the evidence warrants it a narrow invention, must be limited to a narrow claim construction:

> When the specification makes clear that the invention does not include a particular feature, that feature is deemed to be outside the reach of the claims of the patent, even though the language of the claims, read without reference to the specification, might be considered broad enough to encompass the feature in question…. A patentee may also limit the scope of the claims by disclaiming a particular interpretation during prosecution.

Microsoft Corp. v. Multi-Tech Systems, Inc., 357 F.3d 1340, 1347 (Fed. Cir. 2004) (internal citations and quotations omitted). Thus, Sun's patent should be limited, as explained below, to a system that provides an Internet user with human-friendly information about a server on the Internet using the registered domain name of that server.

B. "Domain Name" (Claims 1, 9, 14, 22, 27, and 35)

The table below shows the parties' respective positions for these terms.

Term	NetApp's Construction	Sun's Construction
"domain name"	A third-party-approved name of a website on the Internet, i.e. a registered domain name.	A name that has a numerical IP address associated with it.

The dispute between the parties is whether a "domain name" must be a true third-party-approved name of a website on the Internet, i.e. a registered domain name (NetApp's position) or merely any name that relates to any numerical address (Sun's position).

1. The Specification Supports NetApp's Construction

The specification leaves no doubt that, when the claims refer to a "domain name," they are referring to the standard third-party-approved domain name ubiquitous on the internet (i.e., a registered domain name). Thus, the specification explains that, "[d]omain names are requested by maintainers of the website and are approved by a third party." '106 at 1:56-57. The patent thus specifies that the term "domain name," as used in the '106 patent, refers to registered, third-party-approved domain names. *See also* Almeroth Decl. at 5 (explaining that "[t]he term 'domain name' is understood in the art to refer to an address of a computer on the Internet, where the address is approved by a third party and stored in the Internet's registry of domain names."). Furthermore, several Internet standards documents incorporated by reference in the '106 patent specification confirm that the ordinary meaning of "domain name" is a registered Internet domain name.[2] Indeed, Internet "domain names" must be registered with a third-party Internet Domain Name registry. Otherwise, Internet computers with IP addresses would be impossible to contact on the Internet using a "domain name," because there would be no other means for determining actual IP addresses using the "domain name."

No other meaning of "domain name" would make sense in light of the features disclosed in the patent. For instance, the specification and the claims both contemplate using a third-party to provide the descriptive "server identification information," required by the claims. For example, the service "whois" referenced in the patent, including the claims as a tool for looking up the type of identifying information that might be of interest to the novice user. "Whois" is nothing more than a tool for looking up *registered* domain names to see what person or entity owns them. *See* Almeroth Decl. at 6 ("Both whois and Internet name servers would

[2] *See, e.g.*, '106 patent at 1:27-44 (incorporating by reference Berners-Lee, "Uniform Resource Locators (URL)," RFC 1738 (1994), as well as "The HTTP/1.0 draft document of Feb. 19, 1996 and the HTTP/1.1 draft document of Jan. 19, 1996…"). The HTTP/1.0 and 1.1 Internet Drafts describe the "domain name" included in URLs by defining the "host" component of the URL as a "legal Internet host domain name"; the URL specification, RFC 1738, describes likewise describes the "host" component of URLs as "fully qualified domain name[s] of a network host". In addition, RFC 1738 references Mockapetris, "Domain Names – Concepts and Facilities", RFC 1034 (1987) which states that, "[t]he impetus for the development of the domain system was growth in the Internet…." All these widely-known standards documents make it clear that a domain name is a registered domain name.

have been known to one of ordinary skill in the art as techniques to look up information about a server based on a registered domain name."). What is more, all the diagrams depicting the supposed invention in action are cutouts of Netscape, a program known to perform a DNS look-up for the purpose of locating a website on the Internet based upon domain names—and of course the figures show Netscape doing just that: looking up a webpage on the Internet. *See* '106 at 5:15-32, and Figures 3A-4B as cited in Almeroth Decl. at 6. Thus, all the uses of "domain names" in the patent are specific uses of *registered Internet* "domain-names."

2. The Prosecution History Supports NetApp's Construction

In prosecuting the '106 patent, the patentee himself *admitted* that the term "domain name" should be understood according to its standard meaning: "Domain Name . . . is used in accordance with standard usage in the field." Amendment A, Dec. 22, 1997, ('106 File History) at 6. Consistent with this, the patentee disclaimed any interpretation permitting domain names to be something other than true, registered Internet domain names. Indeed, in distinguishing another limitation—"additional information about the server"—from the prior art, the patentee acknowledged that this "standard usage" requires domain names to be registered:

> "Unlike domain names, this additional information need not be unique *nor approved by a third party*."

Id. Having argued that the term "domain name" had specific properties in order to persuade the examiner of the validity of his claims, the patentee should not now be heard in litigation to argue the very opposite. *See, e.g., Gillespie v. Dywidag Systems International, USA*, 501 F.3d 1285, 1291 (Fed. Cir. 2007) ("The patentee is held to what he declares during the prosecution of his patent.").

3. Extrinsic Evidence Supports NetApp's Construction

Every dictionary definition cited by *either* party confirms that domain names are registered with a third party organization:

- The *Microsoft Press Computer Dictionary*, 3d Ed., (1997) defines the term "domain name" to mean "An address of a network connection that identifies the owner of that address in a hierarchical format: *server.organization.type*. For example, www.whitehouse.gov identifies the Web server at the White House, which is part of the U.S. Government." Identifying the "owner" of the "domain name" only makes sense if the "domain name" is *registered* with

some third party, because without registration there is no ownership of "domain names."

- The same dictionary also defines DNS as "Acronym for **D**omain **N**ame **S**ystem. The system by which hosts on the Internet have both domain name addresses (such as bluestem.prarienet.org) and IP addresses (such as 192.17.3.4). The domain name address is used by human users and is automatically translated into the numerical IP address, which is used by the packet routing software…" Explaining that "domain names" are connected to "hosts on the Internet," this definition confirms NetApp's construction.
- The definition in *Webster's New World Dictionary of Computer Terms*, 8th Ed. (2000), identified by Sun in the Joint Claim Construction and Prehearing Statement, explicitly confirms that term "domain name" refers to an Internet address: "**domain name:** in *the system of domain names* used to identify *individual Internet computers*, a single word or abbreviation that makes up part of a computer's unique name (such as watt.seas.virginia.edu). Reading from left to right, the parts of a domain name go from specific to general…"

All three of these definitions, including most notably the dictionary cited by Sun, define the term "domain name" in the same fundamental way: a "domain name" is a registered domain name, used for communicating with a computer that is on the Internet. As all of the foregoing makes clear, the specification, prosecution history, and all of the extrinsic evidence are aligned with NetApp's construction.

Consistent with this well-known plain and ordinary meaning, the Federal Circuit has consistently used the term "domain name" as a third-party-approved registered Internet domain name. In a variety of cases the Federal Circuit's discussions of domain names are premised on uniqueness of ownership, enforceable property rights, and a centralized system of recording title, all of which are features exclusive to the use of domain names on the Internet, i.e. registered domain names. *See, e.g.*, *In re Oppedahl & Larson LLP*, 373 F.3d 1171, 1176-77 (Fed. Cir. 2004) ("[T]he simple fact that domain names can only be owned by one entity does not of itself make them distinctive or source identifying,"); *Teva Pharmaceuticals USA, Inc. v. Pfizer Inc.*, 405 F.3d 990, 997 (Fed. Cir. 2005) ("United States law 'provides a registrant who has lost a domain name…with a cause of action for an injunction returning the domain name if the registrant can show that she is in compliance with' United States law…"). In explaining URLs, WWW, and DNS, the Federal Circuit confirms the common understanding that domain names are used to access resources on the Internet, a process depending on third-party approval and registration:

> The Internet is a global network connecting millions of computers in more than 100 countries. The World Wide Web, a collection of files, or 'web pages,' containing text, graphics, audio, and video, as well as 'hyperlinks' to other web pages, has become a central part of the Internet. Consumers typically access the web using client software applications known as web browsers that run on their personal computers.
>
> ***Every web page*** is identified by a ***unique Uniform Resource Locator (URL)***. Web pages are stored on 'web sites,' locations on the World Wide Web comprising one or more computers, known as servers. Every web site has a home page, which is identified by a URL and is the first document users see when they first connect to the web site. *Also associated with each web site is a* ***domain name****, usually* ***part of the URL***.
>
> Each web server typically is identified by a unique 32-bit numeric address known as an Internet Protocol address, or IP address. When a user requests a web page by entering a URL into a browser, *the URL is sent to a* ***domain name system (DNS)*** *server*, which uses a look-up table to *translate the* ***domain name*** *in the* ***URL*** into the IP address of a server associated with the web site being accessed. That IP address is returned to the browser, which then uses the address to initiate a communications session with the server that contains the desired web page.

Resonate Inc. v. Alteon Websystems, Inc., 338 F.3d 1360, 1361-62 (Fed. Cir. 2003). Indeed, Congress has also defined domain name to require third party approval: "The term 'domain name' means any alphanumeric designation which is registered with or assigned by any domain name registrar, domain name registry, or other domain name registration authority as part of an electronic address on the Internet." 15 U.S.C. § 1127.

4. Sun's Proposed Construction is Wrong

By contrast, Sun's proposed construction would cause the relevant claim language to read on technologies that long predated the filing of the application for the '106 patent. For instance, static name-to-IP mapping tables were commonly used for communications between computers on a private network. These tables allowed programs such as "Telnet" to provide information about a server computer to a user on another computer. *See* Almeroth Decl. at 9 (explaining how, under Sun's construction, static name-to-IP mapping would include all the elements of representative claim 1 of the '106 patent). If Sun were to win its proposed construction, the prior art "Telnet" program using static name-to-IP mapping would meet every limitation of the claims of the '106 patent.

For these reasons the term "domain name" means "a third-party-approved name of a website on the Internet, i.e. a registered domain name."

C. "Server Identification Data" (Claims 1, 4, 9-10, 14, 17-18, 23-24, 27, 30, 35-36)

The table below shows the parties' respective positions for these terms.

Term	NetApp's Construction	Sun's Construction
"server identification data"	Human-friendly information identifying a specific web server designed not to be intimidating to inexperienced users of the World Wide Web.	Information that uniquely identifies one server from other servers and can be seen by a user.

The dispute between the parties is whether the term "server identification data" must be human-friendly information (NetApp's position) or not (Sun's position). The '106 patent is directed to helping inexperienced users of the internet understand who is providing them with web content they are viewing. Because "'server identification data' has no special meaning in the art," *see* Almeroth Decl. at 9, it is the specification and the file history that must provide the context for construing this term.

1. The Specification Supports NetApp's Construction

The specification explains that displaying "human-friendly server identification information to a user to better indicate the origin of the information," is the goal of this supposed invention. *See* '106 at 1:15-18. The specification explains that this information is needed to help inexperienced users avoid becoming disoriented and intimidated by "surfing" on the WWW:

> WWW users often become *disoriented* while navigating through the WWW…. Often… the domain name is either not descriptive or misdescriptive of the site. For example does www.sun.com refer to the SUN OIL company or to SUN MICROSYSTEMS, Inc.? IP addresses carry no descriptive value whatsoever…. Even if the domain name is unambiguous, its placement in the URL is *intimidating to inexperienced users* of the WWW.

Id. at 1:45-47, 1:60-2:3. Not surprisingly, the only concrete example disclosed in the specification suggests placing the name of the owner of a website in parentheses next to the URL in the location bar of a web browser (or placing the same information at the bottom of the browser window to the right of the URL, or in the browser's bookmark menu alongside the URLs and web page titles). *See* '106 at Fig. 3C; *id.* at Fig. 4B; *id.* at 6:12-21; Almeroth Decl. at 3, 10. The specification further explains, "The server's identification data includes the text string 'SUN MICROSYSTEMS' as descriptive information. The server's identification data is enclosed in parentheses '()' and separated from the URL by a space." '106 at 5:40-44. Thus the entire

intended purpose of the invention is to reduce the disorientation and intimidation of the inexperienced user.

2. The File History Supports NetApp's Construction

The patentee confirmed this narrow purpose when traversing the examiner's rejections, stating that "The invention addresses the problem of identifying and/or describing a server site to a human user by providing additional information about the server to the user…" Amendment A, Dec. 22, 1997, ('106 File History) at 3. The patentee also specifically distinguished other forms of information which were not claimed: "domain name… hypertext links… URLs… web page titles… and bookmarks." *Id*. Moreover, the patentee stated that, "The 'Server Identification Information' term is used throughout the application to mean descriptive information about a server…. This descriptive information is ***not an IP address***…" Amendment A, Dec. 22, 1997, ('106 File History) at 5-6. Therefore, none of these categories of information is claimed. URLs and domain names are specifically cited by the patentee as being confusing, insufficiently descriptive, misdescriptive, or intimidating; and the patentee stated in the specification that "IP addresses carry no descriptive value whatsoever." '106 at 1:63-64. *A fortiori*, nothing *more* difficult to understand or *less* descriptive than these categories could fall within the boundaries set by this disclaimer. In light of these disclaimers, all that remained for the patentee to claim would have been that narrow category of human-friendly information, of which the "Sun Microsystems" parenthetical is one example. This is completely consistent with NetApp's proposed construction and utterly at odds with Sun's.

3. Sun's Proposed Construction is Wrong

It is therefore troubling that Sun's proposed construction would cause the claims to read on the very prior art the patentee disclaimed in the file history, i.e. IP addresses. Indeed, Sun's construction would broadly cover the use of IP addresses and URLs, because IP addresses and URLs uniquely identify a server and they can be seen by a user. *See* Almeroth Decl. at 13. Because of the disclaimers discussed above, Sun's proposed construction runs afoul of the same principle cited in *Gillespie*, 501 F.3d at 1291 (Fed. Cir. 2007): "The patentee is held to what he declares during the prosecution of his patent." Indeed, Sun's construction would broaden Sun's

claims so dramatically that they would appear to include any network device of any kind providing any type of information whatsoever about itself. Such a broadening construction simply cannot be correct in light of the above disclaimer, as well as the stated purpose of the supposed invention. *See Apple Computer, Inc., v. Articulate Systems, Inc.*, 234 F.3d 14, 25 (Fed. Cir. 2000), ("[T]he claim must be interpreted in light of the teachings of the written description and purpose of the invention described therein.").

There can be no doubt that the claims of the '106 patent are limited to this narrow purpose. The '106 patentee referred to his purported invention as "this invention", "the invention", or "the present invention", throughout the patent, no fewer than a dozen times, *e.g.*:

- *Specifically*, ***this invention*** is a new and useful method, apparatus and computer program product for presenting human-friendly server identification information to a user to better indicate the origin of information. '106 at 1:15-18.
- ***The invention*** provides the computer user with information identifying the server that provides access to data (or to services) of interest to the user. *See id.* at Abstract.
- ***The present invention*** overcomes the disadvantages of the above described systems… *Id*. at 2:24-25 (where the "above described systems" are web browsers, Internet domain names, and confusing URLs).

The Federal Circuit has held that this language limits a patent's scope. *See Verizon Services Corp. v. Vonage Holdings Corp.*, 503 F.3d 1295, 1308 (Fed. Cir. 2007), ("When a patent thus describes the features of the '*present invention*' as a whole, this description limits the scope of the invention."). Thus, in light of the repeated assertions throughout the specification and file history stating that the information about a sever is useful to the novice Internet user, "server identification data" must be interpreted to require that the information be in a form that is helpful to the inexperienced user.

For these reasons the term "server identification data" means "human-friendly information identifying a specific web server designed not to be intimidating to inexperienced users of the World Wide Web."

II. U.S. PATENT NO. 5,459,857

A. Background

Sun's 857 patent is directed to a fault tolerant data storage system. A problem cited in the patent is the supposed absence of an inexpensive fault tolerant data storage subsystem

configuration "that minimizes the possibility of a single failure in the storage control unit causing the loss of access to data stored on the disk drives within the data storage subsystem." *See* '857 at 1:47-57

The patentees purported to solve this described problem by connecting two disk array data storage subsystems to each other such that each would send data received from one host processor to the other subsystem. The two disk array data storage subsystems, interconnected by at least one data link, emulate the operation of two storage control units sharing a single extremely reliable physical disk drive. The duplication of data on disks attached to the two operationally independent disk array data storage subsystems is described as transparent to the host. After the host sends data to either disk array data storage subsystem, that disk array data storage subsystem writes that data to its disks and transmits the same data over the data link to the other subsystem to write to its disks. *See generally id.* at 1:60-2:43. Data would thus be stored on two separate storage subsystems, such that the failure of a single subsystem would not lead to the loss of data. The claimed apparatus is depicted in the '857 at Figure 1:

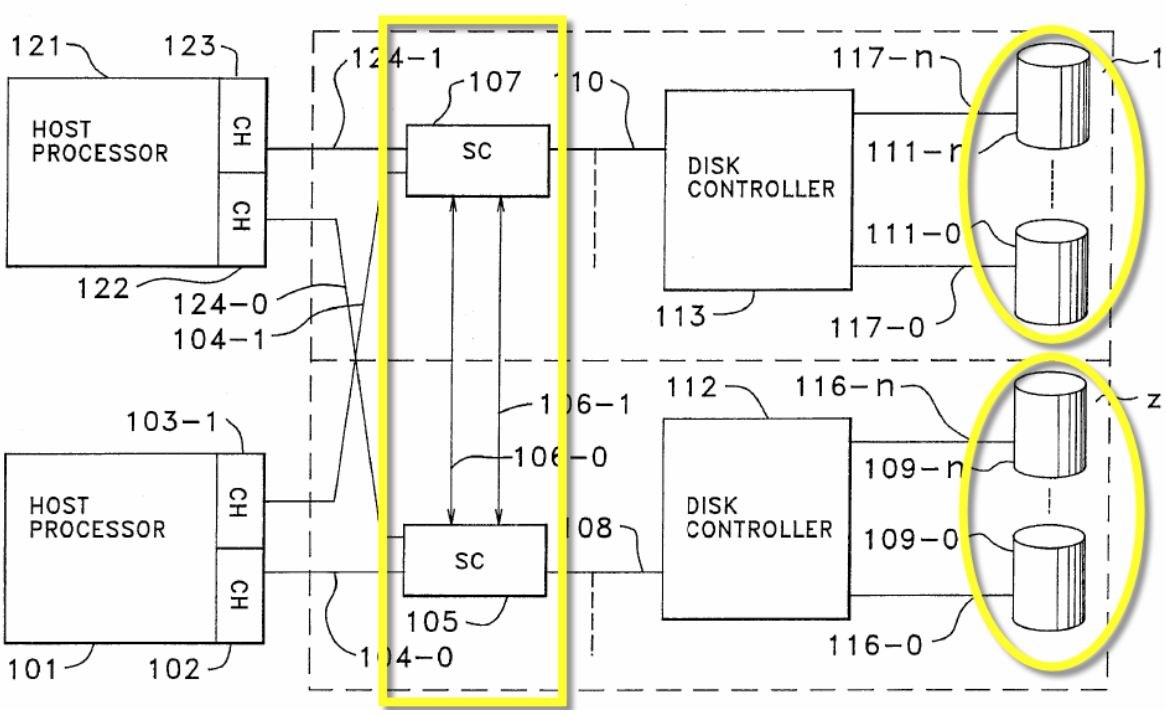

"In response to writing a data record to said one redundancy group" / "Responsive to writing a data record to one of said redundancy groups" / "In response to the receipt of a stream

of data records from said data processor" / "Responsive to the receipt of a stream of data records from said data processor" (Claims 6 and 11)

The table below shows the parties' respective positions for this term.

Term	NetApp's Construction	Sun's Construction
"In response to writing a data record to said one redundancy group" / "Responsive to writing a data record to one of said redundancy groups"	After and in reaction to the writing of a data record to a single redundancy group.	In response to writing the data record to memory associated with the one redundancy group. / Responsive to writing the data record to memory associated with one of the redundancy groups.
"In response to the receipt of a stream of data records from said data processor" / "Responsive to the receipt of a stream of data records from said data processor"	After and in reaction to the receipt of data records from a processor.	Sun contends this phrase does not require construction because the phrase is clear on its face.

Exemplary claim language pertinent to this dispute is as follows:

> *transmitting*, in *response* to writing a data record to said one redundancy group, said *written data record* to the other of said data storage subsystems via said data link to maintain duplicate data records in both said data storage subsystems

This claim language, on its face, requires that the transmission of the "written" data record from one data storage system to another occur only *after* the record is written to a single redundancy group. Likewise, other claim language adopting a matching grammatical structure requires that the selection of data storage space occur only *after* the receipt of data records from a processor.

Why, then, is there a debate worthy of this Court's time? The apparent reason for this is that, for the claim terms referring to "writing a data record to said one redundancy group," Sun believes that, in the preferred embodiment, this "writing" occurs after the step of "transmitting" data to the other storage subsystem. Accordingly, Sun will likely attempt to invoke the oft-repeated claim construction guide that patents are normally construed to cover the preferred embodiment.

Sun's argument fails on two grounds. First, while patents are normally construed to cover the detailed description of the preferred embodiment, not all the claims should logically be expected to do so. *See Helmsderfer v. Bobrick Washroom Equipment, Inc.*, 527 F.3d 1379,

1383 (Fed. Cir. 2008), ("It is often the case that different claims are directed to and cover different disclosed embodiments. The patentee chooses the language and accordingly the scope of his claims."). This principle is particularly applicable here. The patent drafter attempted to claim the purported invention such that it would cover a multitude of different embodiments. Indeed, the '857 patent has 54 different claims, including 16 independent claims. This is more than four times the average number claims in patents of that era and more than eight times the average number of independent claims.[3] Where a patent boasts so many different claims but only one detailed description of the preferred embodiment, it does not take a statistician to recognize that the some claims may cover an embodiment that is not preferred.

Second, it is simply wrong to say that the specification of the '857 contemplates only a single embodiment. Indeed, in the Joint Claim Construction and Prehearing Statement, Sun relies on a passage of the '857 Patent in support of its claim construction position that describes both the disk "writing" step and the "transmitting" step and that requires no particular sequence for the performance of these steps. *See* '857 at 4:39-50. Likewise, the patent never calls out a particular sequence for the "writing" and "transmitting" steps as corresponding to the "preferred embodiment." Nor does it disclaim a particular sequence for performing the "writing" and "transmitting" steps as inefficient. Thus, the fact that one specific implementation of the alleged invention differs from the implementation claimed in two of the 54 claims is irrelevant, because those claims fall within the scope of the general description of the alleged invention.[4]

[3] *See* D. Crouch, *The Rising Size and Complexity of the Patent Document*, Legal Studies Research Paper Series, Research Paper No. 2008-04, Univ. of Mo. School of Law, p. 4 (2008), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1095810.

[4] Sun proposes the construction, "in response to writing the data record to *memory associated with* the one redundancy group." The only substantive change in Sun's construction is "memory associated with." This construction does nothing to clarify the meaning of the term. In the best case it is merely confusing and obfuscating. In the worst case it is intended to impermissibly alter the plain meaning of the claim. If Sun understands correctly that the redundancy group consists of disk drives and that writing to a redundancy group means writing to those disk drives, then its construction is merely confusing—the memory is *part of* the disk drives, not *associated with* them. If on the other hand Sun intends the memory associated with the redundancy group to be something other than the disks themselves, then in addition to being confusing, the construction is simply wrong.

1. The Plain Language Of The Claims Demonstrates That The Required Transmission Of The "Written Record" Occurs Only *After* The Record Is Written To The First Redundancy Group, and That the Required "Selecting" Happens Only *After* the Data Records are Received

As the Federal Circuit has explained repeatedly, "[a] claim must be read in accordance with the precepts of English grammar." *See In re Hyatt*, 708 F.2d 712, 714 (Fed. Cir. 1983); *see also Adang v. Fischhoff*, 286 F.3d 1346, 1352 (Fed. Cir. 2002) ("The court must look at the language of the count as a whole and consider the grammatical structure and syntax."). Applying this basic principle confirms that the asserted claims require that "transmitting" occur only after "writing," and that "selecting" occur only after "receipt."

First, the claim language pertinent to the dispute regarding the sequence of "transmission" and "writing" appears as follows in asserted Claim 6[5]:

> *transmitting*, in *response* to writing a data record to said one redundancy group, said *written data record* to the other of said data storage subsystems via said data link to maintain duplicate data records in both said data storage subsystems

Two aspects of the above quoted claim language make clear that the "transmitting" of the data record occurs only *after* the data record is written to the redundancy group. First, the claim language explains that "transmitting" is done "in response" to the writing of the data record to the redundancy group. Using the term "in response," the plain claim language confirms that the "transmitting" happens after the "writing." Second, the claim language further requires that the "written" data record be the thing that is "transmitted" to the other redundancy group. Using the past-tense form "written" to refer to the data record that is transmitted, necessarily the writing has already occurred before the "transmitting." Indeed, the "transmitting" step must be understood in the context of the "writing" step which by the plain terms has already occurred. *See Pause Technology LLC v. TiVo Inc.*, 419 F.3d 1326, 1331 (Fed. Cir. 2005) ("[P]roper claim construction demands interpretation of the entire claim in context, not a single element in isolation.") (citation omitted).

[5] Nearly identical language appears in asserted Claim 11, and the following analysis is thus equally applicable to Claim 11.

Likewise, the claim language pertinent to the dispute regarding the order of "receipt" and "selecting" appears as follows in asserted Claim 6[6]:

> *selecting*, in *response* to the receipt of a stream of data records from said data processor, available memory space in one of said redundancy groups to store said *received* stream of data records thereon

Furthermore, the structure of this limitation parallels the limitation above pertaining to "writing" and "transmitting," demonstrating a deliberate attempt on the part of the drafter to prepare claims in which pairs of steps occur in a particular sequence. Again, two aspects of the above quoted claim language make clear that the "selecting" of available memory space occurs only *after* the stream of data records is received. First, the claim language explains that "selecting" is done "in response" to the receipt of data records from said data processor. Using the term "in response," the plain claim language confirms that the "selecting" happens after the "receipt." Second, the claim language further requires that the "received" stream of data records be the thing for which the process of "selecting" is initiated. Using the past-tense form "received" to refer to the stream of data records for which storage space is selected, it must be the case that the receipt has already occurred before the "selecting."

2. All 54 Claims Of The '857 Patent Should Not Be Expected To Cover The Most Specific Disclosure Of The Disclosed Embodiment

As noted above, the Federal Circuit recently explained that, logically, every single claim need not cover the preferred embodiment in the patent where there are at least some claims which do so. *See Helmsderfer*, 527 F.3d at 1383 (holding that, properly construed, the claims-in-suit would not cover the preferred embodiment and that there was nothing wrong with this outcome); *Sinorgchem Co., Shandong v. International Trade Commission*, 511 F.3d 1132, 1138 (Fed. Cir. 2007) ("Where, as here, multiple embodiments are disclosed, we have previously interpreted claims to exclude embodiments where those embodiments are inconsistent with unambiguous language in the patent's specification or prosecution history.").

This common sense rule is nowhere more applicable than here, where a patent has been assembled as a diverse array of claims with the apparent goal of scattered patent coverage.

[6] Again, nearly identical language appears in asserted Claim 11

An analysis of the dozens of unasserted claims in the '857 patent confirms that the patentee did specifically claim the most thoroughly described disclosures of the sequence of events referenced in the detailed description, but simply did not choose to do so in every single claim. For instance, Claim 31 covers the transmission of the data record immediately in response to the receipt of the data record from the processor without first requiring that it be written to disk. Indeed, the claim calls for "means… in response to *the receipt* of said data record… for transmitting said data record…" and does not refer to the transmission of the "written" (i.e., past tense) data record. '857 at 20:37-41. A number of other claims include similar limitations. *See, e.g., id.* at Claims 24, 29, ad 33. The patentees obviously understood the distinction between acting in response to "receipt" of a data record, and acting in response to the subsequent storage of that data record. In this litigation, Sun chose not to assert any of the claims requiring action in response to "receipt."

For these reasons the term "in response to" / "responsive to" means "after and in reaction to."

III. U.S. PATENT NO. 5,749,095

A. Background

Sun's '095 patent is directed generally to a technique for improving the performance of a multiprocessing computer system by providing a mechanism for carrying out special "fast write" operations that the processor can complete more quickly than traditional write operations. The accused functionality is isolated to certain specific operations in AMD micro-processors. By way of background, NetApp does not design, develop, nor manufacture microprocessors. To the extent microprocessors are used in NetApp products, they are simply off-the-shelf components purchased from third parties.

The '095 patent is directed to the use of "fast write" operations in *multiprocessing* systems, where issues of "coherency" can arise. At the heart of the alleged invention of the '095 patent is a "system interface" that manages the "fast write" operations and deals with the aforementioned "coherency" issues. These "coherency" issues can arise from the fact that multiprocessing systems may maintain multiple copies of a single piece of shared data – one copy for each processor, for example. More specifically, each processor in a multiprocessing system

typically has associated with it a small memory (a "cache"), which the processor can access rapidly and can use to store frequently needed data. A processor can see data in its own cache, but not the data in the caches of other processors. To illustrate the coherency problem, consider a situation where more than one cache is holding the same piece of data. In this situation, a problem may arise if one processor should choose to modify the copy of the data it is holding in its cache. Indeed, should a processor modify a copy of data only in its own cache, then different processors may no longer have a uniform view of the data: the processor that has modified the data will view the data as having the new value, while the processors that have not modified the data will view the data as having the old value.

At the time the '095 patent was filed, it was long known that to remedy such "coherency" problems arising during certain write operations, coherency activity must be carried out. Specifically, the processor modifying the data must do one of two things. It must either (1) inform the other processors that the copies of the data they are holding are invalid, or (2) provide an updated copy of the data to the other processors. *See* Dubois Decl. at 2.

The '095 patent claims a "fast write" technique. This technique is designed to expedite those operations that are sufficiently independent of other on-going operations that they simply do not implicate immediate coherency issues. Such operations can thus be processed "fast," meaning without the need to await the resolution of coherency processes. Further, a special encoding can be used to inform the system that such "fast write" operations do not need to await coherency processes and they can be processed right away.

As set forth in representative Claim 11, for instance, particular operations in one processer (said processor) can be completed without awaiting completion of coherency operations if the operations are coded for a "fast write":

> wherein said system interface is configured to complete said write operation with respect to said processor prior to completing said coherency operation if said write operation includes a specific predefined encoding, and wherein said system interface is further configured to inhibit completion of said write operation with respect to said processor until completion of said coherency operation if said write operation includes a different encoding

The specification highlights the central role in the claimed invention of the "system interface," depicted below as block 24 in Figure 2 of the '095:

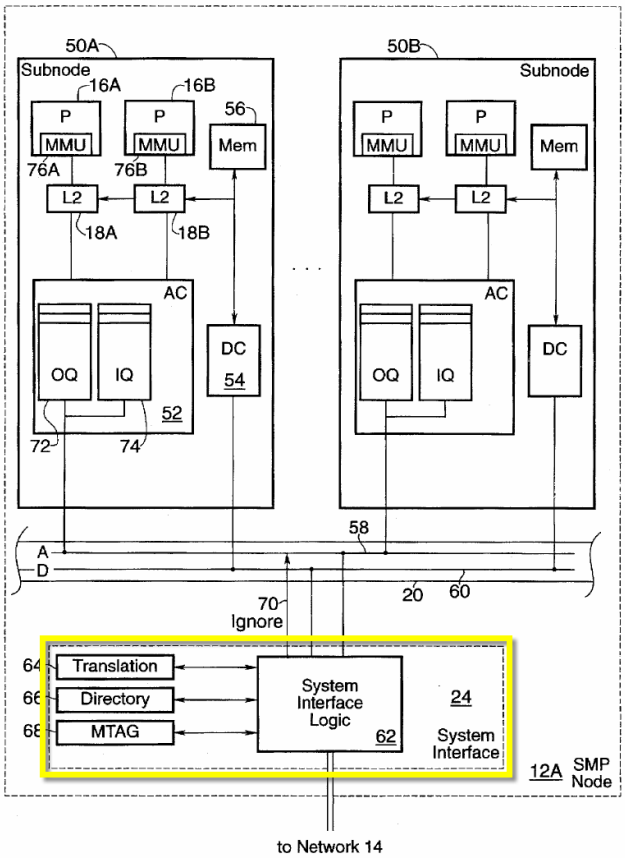

The "system interface" is coupled to the processor and acts as a gateway between the processors and a "network" that the processors in a multiprocessing system use to communicate with each other. As the specification explains, the "present invention" contemplates a "system interface" that carries out a number of core functions pertinent to the execution of write operations, including (1) receiving a write operation, (2) performing coherency operations in response to a write operation, (3) completing "fast write" operations with respect to a processor, and (4) inhibiting the completion of traditional write operations pending the completion of coherency activity. *See* '095 at 5:53-65.

In greater detail, the procedure by which the "system interface" is used to carry out a "fast write" operation is as follows. First, the "system interface" detects a "fast write" operation by its special encoding. *See id.* at 27:15-16. Upon detecting the "fast write" operation, the

"system interface" declines to assert a special signal, called the "ignore signal," which is normally asserted to inhibit further processing of the operation pending the completion of coherency activity. *See id.* at 12:26-43. Instead, the write data is immediately transferred from the initiating processor to the storage queue within the "system interface" *prior* to the completion of the requisite coherency activity. *See id.* at 27:16-25. This storage queue is depicted below as block 94 in Figure 16 of the '095 patent:

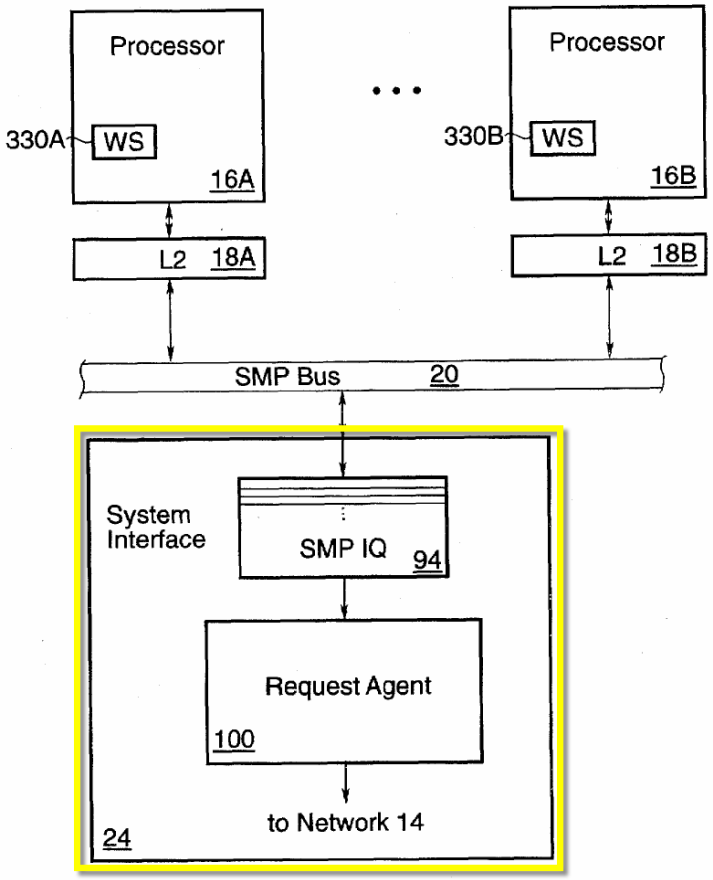

As a result of the write data being transferred to the storage queue in the "system interface," the specification explains that the write operation may "appear to the issuing processor 16 to complete" or, put another way, be "completed in the local node." *See id.* at 27:25-28; *id.* at 27:40-47. Thus, the specification describes the "completion" of write operations in terms of the "system interface."

Moreover, the specification explains the special feature of the "system interface" that allows this type of "completion" to yield tangible performance benefits. Specifically, as compared to a storage buffer on a processor itself, the "system interface" storage queue is much

larger. *See id.* at 29:21-55. Thus, rather than storing a small number of transactions in the processor pending the completion of "coherency operations," which may impede subsequent operations that require the same processor resources, a large number of transactions can be stored in the off-processor "system interface." *See id.* By using the "system interface" in this manner, the specification explains, "[p]rocessor resources are freed more rapidly than if the coherency state is acquired prior to receiving the data from the processor." *Id.* at 27:29-31; *see also id.* at 30:63-31:13 (explaining that when the "ignore signal" is not asserted, data is transferred on the bus connected to the "system interface" so that "processor 16 resources used to store and perform the fast write stream transaction are freed rapidly, allowing the resources to be used for subsequent transactions . . . ."); Dubois Decl. at 3-4.

B. "Completing [a] Write Operation Within [a] Local Processing Node" / "Completing [a] Write Operation With Respect to [a] Processor" (Claims 1, 11, and 17)

The table below shows the parties' respective positions for these terms.

Term	NetApp's Construction	Sun's Construction
"Completing [a] Write Operation Within [a] Local Processing Node"	Transferring the write data from an initiating processor to a system interface.	Sun contends this phrase does not require construction because the phrase is clear on its face. However, if the Court decides the construe this phrase, the phrase should be construed to mean: "data for the write operation (1) is provided to subsequent read operations within the local processing node to the same address as the write operation and (2) is or will be coherent within distributed shared memory.
"Completing [a] Write Operation With Respect to [a] Processor"	Transferring the write data from an initiating processor to a system interface.	Sun contends this phrase does not require construction because the phrase is clear on its face. However, if the Court decides the construe this phrase, the phrase should be construed to mean: "data for the write operation (1) is provided to subsequent read operations by a processor to the same address as the write operation and (2) is or will be coherent within distributed shared memory.

The dispute between the parties is, first, whether construction of these terms is necessary to assist the jury (NetApp's position) or not (Sun's position). Second, the parties

debate what, in the context of the patent, the "finish line" is for a write operation to be considered "complete" "with respect to said processor" or "within said local processing node."

As a threshold matter, claim construction of this term is warranted because Sun has attempted to obscure what constitutes completion of the write operation within the context of this patent and this claim. As explained below, "completion" "with respect to said processor" or "within said local processing node" occurs when the write data is transferred to the system interface. This is because, according to the patent, the key event for such completion to occur is the freeing up of processor resources so it can perform tasks other than the write operation. This happens when the processor off-loads to the "system interface" the data it is attempting to write.

As an initial case in point, in stepping through the "fast write" procedure to which the claims are directed, the specification explains that the "system interface" first recognizes such operations as special operations. *See* '095 at 27:15-16. Then, the same passage explains how the "system interface" may be used to process this "fast write" operation**:**

> Instead of first acquiring a coherency state for the affected coherency unit consistent with performing a write operation and then subsequently transferring the data from the initiating processor, *system interface 24 allows transfer of the data to system interface 24 prior to completing the requisite coherency operation.*

Id. at 27:16-23. In other words, in contrast to other write operations, "fast write" operations allow for the transfer of data to the "system interface" prior to the completion of the "requisite coherency operation." The passage goes on to explain that this causes the processor to view the write operation as being "complete":

> The write operation . . . may thereby *appear to the issuing processor 16 to complete* before the obtaining of the write permission. . . ."

Id. at 27:25-27. In clarifying this, the passage further explains how this corresponds to "completion in the local node":

> [T]he "fast write" write operation is effectively completed outside of the global ordering of computer system 10 since *the operation is completed in the local node prior to acquiring a coherency state* consistent with performing a write operation.

Id. at 27:43-47. Thus, the specification equates the recited completion "with respect to said processor" and completion "within said local processing node" to the transfer of the write data to the "system interface."

This transfer of data to the "system interface" is core to the patent. Indeed, according to the specification, this transfer out of the processor and into the "system interface" provides specific benefits: "Processor resources are freed more rapidly than if the coherency state is acquired prior to receiving the data from the processor." *Id.* at 27:15-31. These benefits are reiterated through the patent:

- The specification explains that "[p]rocessor resources are freed upon transmission of the write operation and corresponding data to the *system interface*, before an appropriate coherency state is acquired by the node containing the processor." *See id.* at 31:36-53.
- The specification explains that the transfer of data upon the data bus used to convey information to the system interface results in "processor 16 resources used to store and perform the fast write stream transaction [being] freed rapidly, allowing the resources to be used for subsequent transactions such as another write stream operation." *See id.* at 30:63-31:13; *see also id.* at Fig. 2 (depicting the referred to data bus as being connected to the system interface).
- The specification explains that "transfer of the data from the initiating processor, [frees] local node resources more quickly than if the same NUMA write transaction were performed using a non-fast write encoding." *See id.* at 28:15-33.

The benefits of the alleged invention of the '095 patent, therefore, are intimately connected to the rapid transfer of data to the "system interface," and the claims should be understood accordingly. *See Apple Computer, Inc.*, 234 F.3d at 25 (Fed. Cir. 2000), ("[T]he claim must be interpreted in light of the teachings of the written description and purpose of the invention described therein.").

Indeed, in addition to the instance described above, other portions of the specification confirm that the recited "completion" steps should be understood in terms of data transfer to the "system interface." For instance, the specification explicitly explains what it means to "complete a write operation with respect to a processor" as follows:

> Since ignore signal 70 is not asserted upon the fast write transaction, the corresponding data is subsequently provided by processor 16 upon data bus 60 (shown in FIG. 2). During step 322, the data is received and stored by system interface 24. *The write operation is thereby complete with respect to the initiating processor* 16.

at 28: 46-52. As yet another example, the Abstract explains that upon detecting a fast write operation, "the data is transferred to the system interface from the processor" and that the "coherency activity employed to acquire the proper coherency state is initiated subsequent to or in

parallel with the receipt of data from the processor." *Id.* at Abstract. The Abstract then explains that, as a result, the "fast write operations are *performed* prior to acquiring write permission to the coherency unit." *Id.* at Abstract. In offering a nearly identical description later in the specification, the '095 patent explains that this results in the processor being freed up to perform other tasks. *Id.* at 5:7-16. Stating that the write operations are "performed," and connecting this to the benefits of the invention, these descriptions again demonstrate that the relevant "completion" in the '095 patent corresponds to the transfer of the write data to the "system interface."

As a final notable example, the specification describes special characteristics of the "system interface," and again connects its use to the recited "completion" steps. Specifically, the '095 patent explains that storage queues on the processors themselves are relatively small, able to hold only a few write transactions at a time. The storage queue within the "system interface," however, is "much larger than the buffers included within processors." *See* 29:33-36. The specification explains the upshot of this:

> Due to the larger number of storage locations within SMP in queue 94, a large number of fast write stream operations may be queued therein. *Since the fast write stream transactions are completed from processors* 16 by storing the transaction into SMP in queue 94 and the corresponding data within output data queue 90, processors 16 may continue with other operations while system interface 24 completes the write stream operations.

Id. at 29:45-52. In other words, rather than storing a small number of transactions in the processor pending the completion of "coherency operations," which may impede subsequent operations that require the same processor resources, a large number of transactions can be stored in the off-processor "system interface." Thus, the specification describes the unique nature of the "system interface" and connects it to the benefits provided by the '095 patent, confirming that to the extent there is an invention in the '095 patent, it resides within the "system interface." Indeed, this passage again specifically equates write operations being "completed from processors" – which is the focus of the claims – with the transfer of write data to the "system interface."

Though Sun's primary position is that these terms are "clear on their face," these phrases, standing alone, simply do not have a well understood meaning in the art. *See* Dubois

Decl. at 7 (explaining that Sun's definition simply describes the way a write operation should be treated during its lifetime in the system, but adds no clarity to the meaning of the claim terms). More importantly, Sun's alternative constructions dodge completely the role of the "system interface" that is so central to the patent. Indeed, as noted, the claims do not simply refer generically to "completing a write operation," but refer to "completing a write operation within said local processing node" and "completing a write operation with respect to said processor." Had the patentee not intended the phrases "within said local processing node" and "with respect to said processor" to give additional meaning to the generic phrase "completing a write operation," the patentee would have simply used the simpler non-specific phrase "completing a write operation." Yet, for the term "completing a write operation within said local processing node," Sun proposes a construction that merely parrots the claim language "within a local processing node," thus ignoring the "system interface" and, even worse, failing to attribute any actual meaning to this part of the claim term. Similarly, for the claim term "completing a write operation with respect to a processor," Sun simply substitutes the phrase "with respect to a processor" with the phrase "by a processor," which again fails to adequately set forth the meaning of the full claim term. However, as explained above, the full claim terms refer to the transfer of write data to the "system interface," which is the precise mechanism that provides the purported benefits of the alleged invention of the '095 patent. The Court should thus not accept Sun's attempt to curtail the importance of the full claim term.

IV. U.S. PATENT NO. 6,873,630

A. Background

Sun's '630 patent is directed generally to network architectures capable of operating at high data transmission rates. Specifically, Sun claims that the '630 patent covers any product that implements the IEEE 802.3 10-Gigabit Ethernet standard for networking. However, NetApp neither designs, develops, nor manufactures the networking equipment that may actually implement this standard, and the accused functionality resides entirely in third-party off-the-shelf components that are ancillary to NetApp storage products. Indeed, Sun's infringement contentions for this patent point to no actual NetApp product specifications, instead referring

largely to IEEE standards documents. To the extent Sun points to an actual NetApp document referring to 10-Gigabit Ethernet, it merely points to a document summarizing the general state of the market for 10-Gigabit Ethernet networking that explains that "10-Gigabit Ethernet is not *about* storage. 10-Gigabit Ethernet is about *IT infrastructure*." *See* David Dale, *The Road to 10-Gigabit Ethernet*, § 6 (2006). Nevertheless Sun accuses NetApp of infringing 53 different claims from this patent.

According to the '630 patent, networked computer systems running multi-media, database, and modeling applications, for example, would welcome the ability to communicate at data rates exceeding 1 gigabit per second. To this end, the '630 patent describes a scheme in which a network communication, rather than being transmitted across a single channel, is divided into multiple logical channels by a network interface for subsequent transmission. According to the patent, the overall data transmission rate, rather than equaling the data rate for a single channel, is then approximately the sum of the data transmission rates for each of the logical channels. *See* '630 patent at 2:36-47; Acampora Decl. at 2-3.

In greater detail, the specification explains the basic mechanism underlying the alleged invention of the '630 patent for an exemplary configuration involving four channels as follows:

In one embodiment of the invention a distributor (e.g., distributor **204** of FIG. 2) accepts a stream of bytes from a MAC module or layer (e.g. a frame), and distributes individual bytes into sub-streams (e.g., mini-frames) in a round-robin fashion. As depicted in the embodiment of FIG. 2, four channels may be implemented with a 10GMII that is four bytes wide; therefore, each time the distributor receives another four bytes, one byte is submitted to each channel. In this manner, an Ethernet frame is divided into four mini-frames for transmission across a different channel.

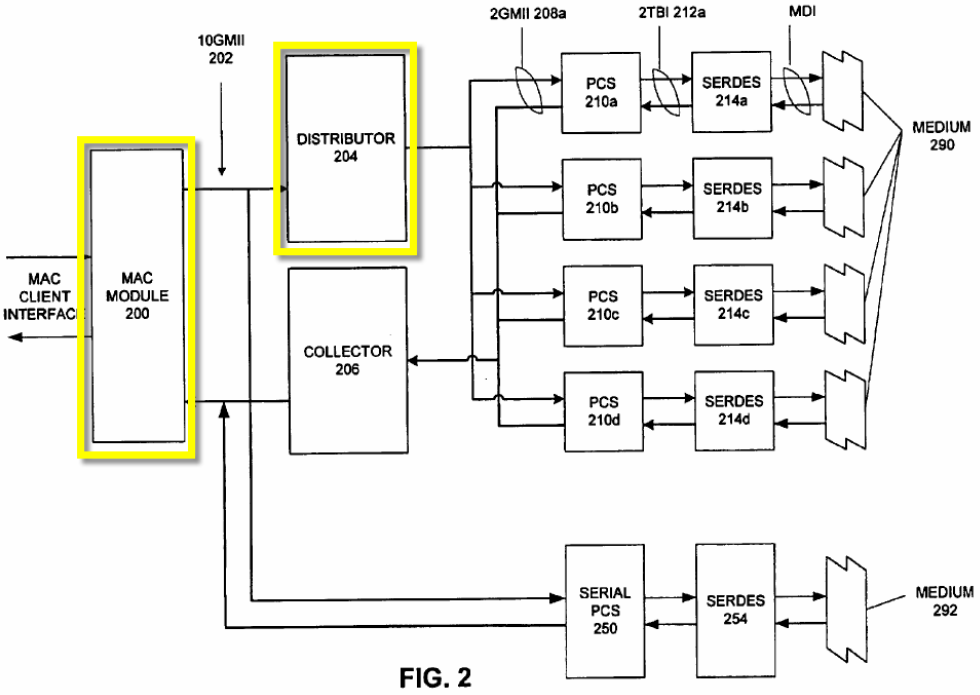

FIG. 2

'630 at 9:24-33. Thus, the '630 patent discloses a scheme in which the communication may be thought of in terms of three levels of granularity. First, and at the highest level, there are the "frames" that are received from a MAC module. These "frames" are then divided into individual "mini-frames" for transmission along individual communication channels. Finally, and at the lowest level of granularity, the "mini-frames" are made up of individual bytes that are ultimately encoded for transmission across the physical medium. *See id.* at 11:40-12:4. By breaking down a single communication, which is normally sent over a single channel, into smaller elements and transmitting them across multiple channels, the '630 patent explains that a higher data transmission rate may be achieved. *See* Acampora Decl. at 2-3.

B. "Portion [of a] Communication" (Claims 3, 5, 12-15, 21-22, 45, 48, 50, 52-53, 73, 76)

The table below shows the parties' respective positions for this term.

Term	NetApp's Construction	Sun's Construction
"Portion [of a] Communication"	The fraction or portion of a frame carried by one channel.	Sun contends this phrase does not require construction because the phrase is clear on its face.

The dispute between the parties is, first, whether construction of this term is necessary to assist the jury (NetApp's position) or not (Sun's position), and, second, whether Sun should be allowed to sidestep the genuine meaning of this term that one of ordinary skill in the art would derive from the specification.

The '630 patent is filled with disclaimers and patent prosecutor fudge-words that try to avoid any committed characterization of what the inventor alleges to have invented. For instance, nothing in the specification forthrightly describes the "invention" itself. Rather, the specification meticulously relies on boilerplate disclaimer paragraphs, and refers only to "embodiments of the invention," cynically describing every aspect of the invention as merely exemplary. However, as explained below, one of skill in the art, after cutting-through the patent prosecutor jargon, would understand that the '630 patent discloses and claims as an invention only a scheme for striping across multiple channels communications received from a MAC module (i.e., "frames"), and the claims should be understood accordingly.

1. A "portion of a communication" is a "fraction or portion of a frame"

Although the patent claims use the vague term "communication," the specification is clear that this must refer to a "frame," where a "frame" is the communication unit emitted from or received by a MAC sublayer in the OSI reference model. This understanding of the term "frame" is well known to those of skill in the art. *See* Acampora Decl. at 4. And, it is explicitly confirmed in the specification of the '630 patent also. *See* '630 at 5:48-51 (defining "frame" to refer to the "unit of information received from or sent to a MAC layer from a physical layer device."). As such, the specification is clear that when the claims refer to a "portion of a communication," they must, at the very least, be referring to a "fraction or portion of a frame." The term "portion of a communication" appears in the specification only once. *See id.* at 4:55-61.

There, the specification explains that a "distributor" is responsible for "disseminat[ing] *portions of the communications* across the multiple logical channels." *See id.* But Sun has agreed that a "distributor" is a module that divides "across multiple logical channels an Ethernet frame received from a MAC module." *See* Supplemental Joint Claim Construction and Prehearing Statement, Exh. Q [Document No. 74-8]. Thus, Sun has in fact already agreed that the only time the specification refers to a "portion of a communication," it is in fact referring to a fraction or portion of a "frame."

Sun's agreement to this is entirely consistent with the specification, because the only embodiments disclosed in the specification in any meaningful way are those in which "frames" – received from a MAC sublayer – are divided into the "portions of a communication" that the claims describe as being transmitted along logical channels:

- The specification defines the "mini-*frame*," as a "fraction or *portion* of a *frame* this is sent across one of multiple channels."[7] *Id.* at 5:51-53. The term "mini-frame" is used 45 times throughout the text of the specification to describe the alleged invention, demonstrating the patent's focus on the use of portions of "frames" for striping communications across communication channels.
- Figure 2 of the specification, which corresponds to the only disclosed embodiment, depicts a "distributor" dividing communications from a MAC module (*i.e.*, frames) among multiple logical channels. *See id.* at Figure 2; *see also id.* at 5:48-51 (defining the "frame" as the unit of information "received from or sent to a MAC layer"); *see also* Acampora Decl. at 5 (offering the same explanation for Figure 1 of the '630 patent).
- As NetApp's expert, Anthony Acampora, explains in his declaration, Figures 3A and 3B describe the methods for transmitting and receiving a "packet," which the specification explains should be understood to be synonymous with a "frame" in the context of the '630 patent. *Id.* at 5:48-51. The figures describe these processes in terms of "mini-*frames*." *See* Acampora Decl. at 4.
- Figures 5A-5D depict the conversion of "frames . . . into multiple mini-frames for transmission across separate channels in accordance with one embodiment of the present invention."

[7] Notably, when the specification uses the term "portion" by itself, it is only in connection with the pieces of a frame or packet that are sent across a single logical channel. *See* '630 at 2:48-51; *id.* at 5:51-53; *id.* at 12:19-21. That the specification consistently uses "portion" in connection with the pieces of frame or packet sent across a single channel demonstrates that the claims also use the term in the same way. *See BellAtlantic Network Services v. Covad Communications. Group*, Inc., 262 F.3d 1258, 1271 (Fed. Cir. 2001) ("[W]hen a patentee uses a claim term throughout the entire patent specification, in a manner consistent with only a single meaning, he has defined that term 'by implication.'").

- The specification explains that the "portions of a single packet" delivered by a single channel are received by a "collector." *See id.* at 12:14-20. But the parties have already agreed that a "collector" is an entity that "reassembles an Ethernet frame to be provided to a MAC layer," demonstrating that the "portion of a single packet" must be referring to a portion of a "frame." *See* Supplemental Joint Claim Construction and Prehearing Statement, Exh. Q [Document No. 74-8].

Where, as here, a patent is singularly focused on achieving a goal that is expressly limited to a specific means, the claims should be so understood. *See On Demand Machine Corp. v. Ingram Industries, Inc.*, 442 F.3d 1331, 1340 (Fed. Cir. 2006) (explaining that the "claims cannot be of broader scope than the invention that is set forth in the specification."). This is, of course, *not* to say that patents with broader inventions should normally be limited to their embodiments; they should not.

The '630 patent explains the shortcomings of an approach where something other than "fractions or portions of frame" are used as the "portions of the communication," and "contrast[s]" this with "one or more embodiments of the invention." Specifically, in connection with its description of Figure 2, the specification explains that an approach where something other than "fractions or portions of frame" are used as the "portions of the communication" – apparently used in prior art 802.3 link aggregation – would require "almost all of the network interface hardware used for present Ethernet implementations [to be] duplicated," while the invention of the '630 patent requires the duplication of only physical layer resources. *See id.* at 8:24-43. No additional details or explanation are offered as to how striping could be achieved based on something other than the use of "portions of communications" that are "fractions or portions of a frame." "Where the general summary or description of the invention describes a feature of the invention . . . and criticizes other products . . . that lack that same feature, this operates as a clear disavowal of these other products." *See AstraZeneca AB, Aktiebolaget Hassle, KBI-E, Inc. v. Mutual Pharmaceutical Co. Inc.*, 384 F.3d 1333, 1339-40 (Fed. Cir. 2004). Thus, a "portion of a communication" must be a "portion of a frame."

The claims of the '630 patent that refer to a "portion of a communication" confirm this. Certain of these claims refer expressly to a "communication" as being received or sent to a "media access control module," leaving no doubt that these claims use the term "portion of a

communication" to refer to a "fraction or portion of a frame." If this were not so, then the "communication" would not be received from a "media access control module." The remaining claims that refer to "portion of a communication" also include structure confirming that the "portion of a communication" is in fact a "fraction or portion of a frame." For instance, claim 15 refers to the process of "collecting" a "portion." The "collecting" must be performed by the "collector," which the parties have agreed refers to a "module that reassembles an Ethernet frame to be provided to a MAC layer," confirming that the "portion" must be a "fraction or portion of a frame." *See* Supplemental Joint Claim Construction and Prehearing Statement, Exh. Q [Document No. 74-8].

There can be no dispute that a "portion of a communication" is transmitted along a single channel. Indeed, none of the claims describe a "portion of a communication" as being transmitted across multiple channels. Rather, the "portions of a communication" are always linked to a single channel. For instance, Claim 3 calls for "sending a *first portion* of said communication on a *first channel*," and a "*second portion* of said communication on a *second channel*." *See id.* at 15:21-29. Every embodiment set forth in the specification corresponds to this understanding. For instance, in defining the term "mini-frame," which is used throughout the specification to describe the invention, the specification explains that it is the "portion of the frame carried by *one channel*." *See id.* at 2:48-49; *see also* 5:51-53 (a "mini-frame" or "mini-packet" is sent across "*one* of multiple channels"). Thus, when describing an embodiment consisting of four channels, the specification explains that the "mini-frame" should consist of "approximately one fourth of the original frame," such that there would then be one "mini-frame" for each channel. *See id.* at 9:64-67; *see also id* at 9:48-50 (explaining that the "mini-frames" are the entities that the "collector" receives on "all channels"). Likewise, in describing the process of receiving a packet at a collector, the specification explains that "each channel" will deliver "portions of a single packet," confirming once again that a "portion" is an entity that corresponds to a single channel. *Id.* at 12:14-20.

Though Sun contends that this vague term is "clear on its face," the term actually has no commonly understood meaning to one of skill in the art. *See* Acampora Decl. at 3. Taken

"on its face," as Sun suggests, the term is indistinguishable from the similarly nebulous claim term "element of a communication," which is discussed immediately below. Yet the claims rely heavily on the term "portion of a communication," making a definitive understanding of its meaning critical. And, for the foregoing reasons, the specification and claims make clear that a "portion of a communication" is "the fraction or portion of a frame carried by one channel."

C. "Element [of a] Communication" / "Element [of a] Portion" / "Elements" / " Element of a Portion" (Claims 3, 5, 8, 13-15, 21-23, 45-48, 50, 52, 73, 76, 79, 89, 113, 114, 116, 117)

The table below shows the parties' respective positions for this term.

Term	NetApp's Construction	Sun's Construction
"Element [of a] Communication"	A portion (e.g., a byte) of a mini-frame that is individually encoded for transmission across one of a plurality of logical channels, where a mini-frame is a fraction or portion of a communication received from or sent to a media access control layer from a physical layer device and that is carried by one channel.	Sun contends this phrase does not require construction because the phrase is clear on its face.

The dispute between the parties is, first, whether construction of this term is necessary to assist the jury (NetApp's position) or not (Sun's position), and, second, whether Sun can again dodge the clear meaning of this term that one of skill in the art would understand from the specification.

1. An "Element" is a Part of a "Mini-frame"

The claims draw a distinction between an "element of a communication" and a "portion of a communication." While some claims refer to an "element of a communication portion," demonstrating that an "element" is a part of a "portion," *see, e.g.*, *id.* at 17:9-14; *id.* at 20:50-55, other claims describe the process of "distributing elements of said communication into multiple portions" (i.e. each "portion" is formed from the "elements" that were distributed to it), again demonstrating that an "element" is a component of a "portion" *See, e.g.*, *id.* at 15:19-20; 20:46-47. Every claim that refers to an "element" maintains this distinction, demonstrating conclusively that an "element of a communication" is a part of a "portion of a communication."

Because an "element" is a part of a "portion of communication," and a "portion of a communication" is a "mini-frame," it follows that an "element of a communication" must be part of a "mini-frame." Indeed, the specification explains repeatedly that a "mini-frame" is just the "portion of the frame carried by one channel," *see id.* at 2:49; *id.* at 5:51-53, which, as explained above, is substantively equivalent to how the term "portion of a communication" should be understood. *See supra* Part IV.B.1. When the specification explicitly refers to "elements," it does so in a way that confirms this aspect of the construction. For instance, in describing the procedure for allocating "elements" to channels, the specification explains that these "elements" are allocated to the logical channels on a "round-robin basis" such that each channel ultimately carries one "mini-frame" or "mini-packet." *See* '630 at 6:64-7:6. Because the "elements" are the fundamental units that are allocated to the channels to form the "mini-frames," they must be portions of "mini-frames." Likewise, when the specification describes how Ethernet frames are encoded for transmission along the communication channels, it explains that it is the individual "elements" that are encoded for transmission. *See id.* at 7:40-53. It is further explained that the encoding may be accomplished using the well known 8B/10B procedure for the coding of individual bytes, thus reflecting a situation where the "element" is one byte in size. *See id.* It is then explained that a "PCS module" at "a receiving entity" "decodes the mini-frame" to "provide[] the recaptured bytes to a collector." *See id.* Because it is the "recaptured bytes" that are extracted when a "mini-frame" is decoded, it must be that an "element" is a portion of a "mini-frame." The specification contains no other direct references to the term "element," and contains nothing to suggest that an "element" could be something other than a "portion of a mini-frame." Given the specification's consistent use of the term "element," it must be that the term "element" corresponds to a "portion of a mini-frame."

2. An "Element" is Transmitted Across a Single Channel

As explained above, the claims and specification confirm that a "portion of a communication" is in fact transmitted across only one of a plurality of communication channels. *See supra* Part IV.B.1. Likewise, an "element of a communication" is a part of a "portion of a communication." Because an "element" is only a sub-part of a larger entity that is transmitted

across only a single channel, it must, as a matter of pure logic, be the case that the "element" is also transmitted across only a single channel.

3. An "Element" is Individually Encoded For Transmission

The communication must ultimately be encoded for transmission, and the specification discloses nothing other than the "elements" being encoded. For instance, all of the embodiments depicted in the specification include at least one Physical Coding Sublayer ("PCS"). *See, e.g.*, Figure 2. This is the only structure disclosed in the specification that carries out any type of encoding, and the parties have agreed that it does its encoding at the "physical layer," which is the layer actually associated with the physical transmission medium. *See* Supplemental Joint Claim Construction and Prehearing Statement, Exh. Q [Document No. 74-8]. Each of the PCS modules responsible for carrying out the PCS functions are depicted as being linked to a single transmission channel and receiving input from a "distribution module," which the parties have already agreed is responsible for "divid[ing] across multiple logical channels an Ethernet frame received from a MAC module." *See id.*; '630 at Figure 2. Because the entity responsible for encoding receives its inputs from the entity responsible for dividing frames into "elements," and because the coding module is directly associated with the ultimate transmission medium, it must be the case that the "elements" are encoded for transmission.

This is confirmed elsewhere in the specification. For instance, the specification explains that the "PCS modules . . . perform coding of Ethernet frame *elements* . . . ." *See* '630 at 7:41-43. Similarly, the specification repeatedly and consistently points to bytes (or smaller units) as the unit of a communication that is encoded for transmission:

- The specification explains that "the remainder of the packet is received by the distributor, distributed *one byte at a time* (in round robin fashion) to each channel, encoded, and transmitted." *See id.* at 11: 60-62.
- The specification explains that "the distributor sends the first *byte* of each channel's mini-frame to a PCS for encoding." *See id.* at 11:50-51.
- The specification explains that the PCS module may perform 8B/10B or 4B/5B coding schemes to encode or decode data. *See id.* at 7:41-53; *id.* at 10:8-19. Such schemes operate on a byte or half-a-byte, respectively.
- The specification explains that "the last *byte* of each mini-frame and the last *byte* of the packet are recognized by their distinctive codes," demonstrating that it is the bytes that have been encoded. *See id.* at 12:35-36.

A byte, carrying only 8-bits of information, is too small to correspond to either a "frame" or a "mini-frame," and must thus correspond to the most granular information unit used in the claims, which is the "element." *See* Acampora Decl. at 9-10. Thus, the specification supports the notion that it is the "elements of a communication" that are encoded for transmission. For the above reasons, an "element of a communication" is "a portion (e.g., a byte) of a 'mini-frame' that is individually encoded for transmission across one of a plurality of logical channels," where a "mini-frame" is "a fraction or portion of a communication received from or sent to a media access control layer from a physical layer device and that is carried by one channel."

Dated: July 7, 2008

/s/ Edward R. Reines

Matthew D. Powers
Edward R. Reines
Jeffrey G. Homrig
Jill J. Ho
WEIL, GOTSHAL & MANGES LLP
201 Redwood Shores Parkway
Redwood Shores, CA 94065
Telephone: (650) 802-3000
Facsimile: (650) 802-3100

Attorneys for Plaintiff-Counterclaim Defendant NetApp, Inc.